WEST Search History

DATE: Monday, December 15, 2003

Set Name side by side	Query	Hit Count	Set Name result set
	DB=USPT; PLUR=NO; OP=OR		
L43	(l39 or l40 or l41 or l42) and (b adj1 tree)	47	L43
L42	L38 and (column$ same (boolean adj1 operator))	5	L42
L41	L38 and (column$ same (greater and less and equal))	44	L41
L40	L38 and ((search$ or quer$) same greater and less and equal)	88	L40
L39	L38 and ((search$ or quer$) same boolean)	56	L39
L38	(l36 or l37) and (index$ same database same column$)	362	L38
L37	((707/100)!.CCLS.)	1342	L37
L36	((707/1 |707/2 |707/3)!.CCLS.)	4094	L36
L35	L34 and (column same row$ same (text or word$))	13	L35
L34	L33 and (column$ and table$)	38	L34
L33	L32 and index$	41	L33
L32	L31 and (user adj1 interfac$)	46	L32
L31	L1 and (database adj1 management)	157	L31
L30	L29 and (column$ same signature)	3	L30
L29	index$.ti.	4009	L29
L28	L17 and (column same signature)	0	L28
L27	L17 and (column same multimedia)	1	L27
L26	L17 and pars$	13	L26
L25	L17 and ((invert$ adj1 index) same pars$)	2	L25
L24	L23 and pars$	0	L24
L23	6349308.pn.	1	L23
L22	L20 and table	7	L22
L21	L20 and table	7	L21
L20	L17 and (invert$ adj1 index)	7	L20
L19	L17 and ((search$ or quer$) same (greater or less or equal))	31	L19
L18	L17 and (column$ same boolean)	7	L18
L17	L1 and index$.ti.	53	L17
L16	L14 and index$.ti.	26	L16
L15	L1 and (greater adj1 than)	0	L15
L14	L13 and less	111	L14
L13	L1 and greater	159	L13
L12	L11 and (search$ or quer$)	16	L12
L11	L9 and L8	16	L11

09/928,894

L10	L1 and (equal$ adj1 condition$)	2	L10
L9	L1 and (boolean adj1 operator$)	21	L9
L8	L1 and (index$ same column$ same database$)	126	L8
L7	L1 and (index$ same column$)	162	L7
L6	L1 and (column$ same database)	187	L6
L5	(unstructured adj1 column$)	2	L5
L4	(unstructured near column$)	2	L4
L3	((unstructured near column$) same database)	0	L3
L2	L1 and (unstructured near column$)	0	L2
L1	(5729730 5893087 6151604 6163775 5983232 6463433 6018734 5446882 5544358 5754782 4961139 5682524 5809497 6182121 6236994 5848241 6006215 6021410 6076087 5664109 5706452 5768578 5826085 5832450 5918229 6279006 6304848 6311191 6442567 6446061 6567812 6345126 5799310 6047291 6078925 5727195 5742809 5995957 5511190 5537589 5724577 5732258 5781896 5819256 5974408 6006219 6064999 6173278 6272495 6178419).pn. (6226641 6226641 6167397 5498445 5974420 6292803 6424358 6523028 6105022 6226632 6226632 5960426 6085189 6128610 5257365 5404512 5404513 5442784 5467471 5706495 5758145 5809495 5890147 5966735 6119128 6134546 6182079 6253196 6338056 6366934 6389413 6421656 6424980 6434551 6507846 6519597 6535875 6584459 6243713 6374267 6411969 6636628 5963933 5548758 5706494 5864842 6003033 4506326 5261093 5276870).pn. (5504886 5548755 5548754 5590324 5701460 5701453 5752016 5761494 5787411 5884299 5905985 5926818 5937402 5940818 5950196 5960438 6098075 6105017 6112197 6144983 6236986 6418451 6519603 5905498 6078924 4884218 5317731 5717913 5835856 5881235 5892843 5938690 5987449 6085105 6088594 6587830 5424732 5893104 6035298 6208987 4991087 5379419 5625813 5701498 5701467 5727197 5829004 5852821 5950191 5963965).pn. (5978792 5999940 6049804 6061503 6061503 6094649 6212532 6223269 6212532 6223269 6345288 6349308 6421677 6118760 5361353 5778373 6052686 6052693 6073130 6106561 5454102 5506974 5864848 4334275 4417321 4510567 4587633 4868735 5301332 5481707 5506968 5845303 5936240 5983214 5995961 6189113 6195353 6377965 6473893 6519261 6587852 5860070 5594898 5553218 5604899 5617567 5652882 5675779 5787410 5812998).pn. (5826259 5870747 5918225 6169983 6223182 6223182 5404510 5465352 5644763 5848408 5913206 5930800 5930785 5960428 6009432 6105018 6154748 6272487 4780810 5414834 5524238 5666526 5752025 5761653 5819251 5822748 5842212 5842197 5870765 5870759 5873091 5875445 5884307 5884325 5903888 5913207 5926816 5926813 5950186 5956705 5960423 5974418 5983215 6026406 6047285 6061678 6061690 6061690 6061678 6081805).pn. (6081800 6105033 6105025 6115722 6122644 6134558 6144957 6161105 6167399 6185557 6189010 6199062 6223171 6223171 6243709 6243718 6263338 6266673 6304876 6351742 6356888 6360213 6366902	291	L1

6370522 6374232 6374235 6381607 6397204 6397219 6401089
6405187 6421662 6427123 6427143 6427145 6438562 6442539
6446063 6457014 6466940 6487549 6490578 6502089 6505199
6507835 6513028 6564204 6567816 6598038 6606626).pn.

END OF SEARCH HISTORY

Terms used
indexing and **database** and **column** and **rows** and **boolean** and **query** and **b tree**

Found **16,573** of **124,998**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown

Relevance scale

1 Improved query performance with variant indexes

Patrick O'Neil, Dallan Quass

June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data**, Volume 26 Issue 2

Full text available: pdf(1.54 MB) Additional Information: full citation, abstract, references, citings, index terms

The read-mostly environment of data warehousing makes it possible to use more complex indexes to speed up queries than in situations where concurrent updates are present. The current paper presents a short review of current indexing technology, including row-set representation by Bitmaps, and then introduces two approaches we call Bit-Sliced indexing and Projection indexing. A Projection index materializes all values of a column in RID order, and a Bit-Sliced index essentially takes an orth ...

2 Query processing and optimization in Oracle Rdb

Gennady Antoshenkov, Mohamed Ziauddin

December 1996 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 5 Issue 4

Full text available: pdf(92.62 KB) Additional Information: full citation, abstract, index terms

This paper contains an overview of the technology used in the query processing and optimization component of Oracle Rdb, a relational database management system originally developed by Digital Equipment Corporation and now under development by Oracle Corporation. Oracle Rdb is a production system that supports the most demanding database applications, runs on multiple platforms and in a variety of environments.

Keywords: Dynamic optimization, Optimizer, Query transformation, Relational database, Sampling

3 Bit-sliced index arithmetic

Denis Rinfret, Patrick O'Neil, Elizabeth O'Neil

May 2001 **ACM SIGMOD Record , Proceedings of the 2001 ACM SIGMOD international conference on Management of data**, Volume 30 Issue 2

Full text available: pdf(182.64 KB) Additional Information: full citation, abstract, references, citings, index terms

The bit-sliced index (*BSI*) was originally defined in [ONQ97]. The current paper introduces

09/928,894

h be c g e cf c

the concept of BSI arithmetic. For any two BSI's X and Y on a table T, we show how to efficiently generate new BSI's Z, V, and W, such that Z = X + Y, V = X - Y, and W = MIN(X, Y); this means that if a row r in T has a value x represented in BSI X and a value y in BSI Y, the value for r in BSI Z will be x + y, the value in V will be x - y and the value in W will be MIN(x, y). Since a bitmap repre ...

4 AutoAdmin "what-if" index analysis utility
Surajit Chaudhuri, Vivek Narasayya
June 1998 **ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data**, Volume 27 Issue 2
Full text available: pdf(1.52 MB) Additional Information: full citation, abstract, references, citings, index terms

As databases get widely deployed, it becomes increasingly important to reduce the overhead of database administration. An important aspect of data administration that critically influences performance is the ability to select indexes for a database. In order to decide the right indexes for a database, it is crucial for the database administrator (DBA) to be able to perform a quantitative analysis of the existing indexes. Furthermore, the DBA should have the ability to propo ...

5 On supporting containment queries in relational database management systems
Chun Zhang, Jeffrey Naughton, David DeWitt, Qiong Luo, Guy Lohman
May 2001 **ACM SIGMOD Record , Proceedings of the 2001 ACM SIGMOD international conference on Management of data**, Volume 30 Issue 2
Full text available: pdf(223.70 KB) Additional Information: full citation, abstract, references, citings, index terms

Virtually all proposals for querying XML include a class of query we term "containment queries". It is also clear that in the foreseeable future, a substantial amount of XML data will be stored in relational database systems. This raises the question of how to support these containment queries. The inverted list technology that underlies much of Information Retrieval is well-suited to these queries, but should we implement this technology (a) in a separate loosely-coupled IR engin ...

6 Query processing: Factorizing complex predicates in queries to exploit indexes
Surajit Chaudhuri, Prasanna Ganesan, Sunita Sarawagi
June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on on Management of data**
Full text available: pdf(240.56 KB) Additional Information: full citation, abstract, references, index terms

Decision-support applications generate queries with complex predicates. We show how the *factorization* of complex query expressions exposes significant opportunities for exploiting available indexes. We also present a novel idea of relaxing predicates in a complex condition to create possibilities for factoring. Our algorithms are designed for easy integration with existing query optimizers and support multiple optimization levels, providing different trade-offs between plan complexity and ...

7 Query Optimization in Database Systems
Matthias Jarke, Jurgen Koch
June 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 2
Full text available: pdf(2.84 MB) Additional Information: full citation, references, citings, index terms

8 XPS a database server for data warehousing
Andreas Weininger
November 2001 **Proceedings of the 4th ACM international workshop on Data**

h be c g e cf c

warehousing and OLAP

Full text available: pdf(621.63 KB) Additional Information: full citation, abstract, citings

A database server used for implementing a data warehouse must support other features than a database server used for OLTP. Therefore, in this paper we will look specifically at features necessary for efficiently processing queries on a database with a star schema model, a database scheme which is used very often in data warehousing. We will especially analyze the features provided for this by the IBM Extended Parallel Server (XPS). There are special star join methods like the Push-Down Hash Semi ...

9 Reusing invariants: a new strategy for correlated queries

Jun Rao, Kenneth A. Ross

June 1998 **ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data**, Volume 27 Issue 2

Full text available: pdf(1.55 MB) Additional Information: full citation, abstract, references, citings, index terms

Correlated queries are very common and important in decision support systems. Traditional nested iteration evaluation methods for such queries can be very time consuming. When they apply, query rewriting techniques have been shown to be much more efficient. But query rewriting is not always possible. When query rewriting does not apply, can we do something better than the traditional nested iteration methods? In this paper, we propose a new invariant technique to evaluate correlated queries ...

10 Strategies for processing ad hoc queries on large data warehouses

Kurt Stockinger, Kesheng Wu, Arie Shoshani

November 2002 **Proceedings of the 5th ACM international workshop on Data Warehousing and OLAP**

Full text available: pdf(245.31 KB) Additional Information: full citation, abstract, references, index terms

As data warehousing applications grow in size, existing data organizations and access strategies, such as relational tables and B-tree indexes, are becoming increasingly ineffective. The two primary reasons for this are that these datasets involve many attributes and the queries on the data usually involve conditions on small subsets of the attributes. Two strategies are known to address these difficulties well, namely vertical partitioning and bitmap indexes. In this paper, we summarize our exp ...

11 Data structures for efficient broker implementation

Anthony Tomasic, Luis Gravano, Calvin Lue, Peter Schwarz, Laura Haas

July 1997 **ACM Transactions on Information Systems (TOIS)**, Volume 15 Issue 3

Full text available: pdf(316.45 KB) Additional Information: full citation, abstract, references, citings, index terms, review

With the profusion of text databases on the Internet, it is becoming increasingly hard to find the most useful databases for a given query. To attack this problem, several existing and proposed systems employ brokers to direct user queries, using a local database of summary information about the available databases. This summary information must effectively distinguish relevant databases and must be compact while allowing efficient access. We offer evidence that one broker, GlOSS

Keywords: GlOSS, broker architecture, broker performance, distributed information, grid files, partitioned hashing

12 Extensions to Starburst: objects, types, functions, and rules

Guy M. Lohman, Bruce Lindsay, Hamid Pirahesh, K. Bernhard Schiefer

October 1991 **Communications of the ACM**, Volume 34 Issue 10

Full text available: pdf(5.21 MB) Additional Information: full citation, references, citings, index terms

Keywords: Extended relational database management systems, Starburst, extensible database management systems

13 Performance comparison of property map and bitmap indexing

Ashima Gupta, Karen C. Davis, Jennifer Grommon-Litton

November 2002 **Proceedings of the 5th ACM international workshop on Data Warehousing and OLAP**

Full text available: pdf(250.60 KB) Additional Information: full citation, abstract, references, index terms

A data warehouse is a collection of data from different sources that supports analytical querying. A Bitmap Index (BI) allows fast access to individual attribute values that are needed to answer a query by representing the values of an attribute for all tuples separately, as bit strings. A Property Map (PMap) is a multidimensional indexing technique that pre-computes attribute expressions, called properties, for each tuple and stores the results as bit strings [DD97, LD02]. This paper compares t ...

Keywords: bitmap index, data warehouse, performance study

14 Similarity queries I: Robust and efficient fuzzy match for online data cleaning

Surajit Chaudhuri, Kris Ganjam, Venkatesh Ganti, Rajeev Motwani

June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on on Management of data**

Full text available: pdf(271.47 KB) Additional Information: full citation, abstract, references, index terms

To ensure high data quality, data warehouses must validate and cleanse incoming data tuples from external sources. In many situations, clean tuples must match acceptable tuples in *reference tables*. For example, product name and description fields in a sales record from a distributor must match the pre-recorded name and description fields in a product reference relation.A significant challenge in such a scenario is to implement an efficient and accurate fuzzy match operation that can effec ...

15 Special issue on prototypes of deductive database systems: The aditi deductive database system

Jayen Vaghani, Kotagiri Ramamohanarao, David B. Kemp, Zoltan Somogyi, Peter J. Stuckey, Tim S. Leask, James Harland

April 1994 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 3 Issue 2

Full text available: pdf(2.67 MB) Additional Information: full citation, abstract, references, citings

Deductive databases generalize relational databases by providing support for recursive views and non-atomic data. Aditi is a deductive system based on the client-server model; it is inherently multi-user and capable of exploiting parallelism on shared-memory multiprocessors. The back-end uses relational technology for efficiency in the management of disk-based data and uses optimization algorithms especially developed for the bottom-up evaluation of logical queries involving recursion. The front ...

Keywords: implementation, logic, multi-user, parallelism, relational database

16 Heuristic optimization of OLAP queries in multidimensionally hierarchically clustered databases

Dimitri Theodoratos, Aris Tsois
November 2001 **Proceedings of the 4th ACM international workshop on Data warehousing and OLAP**
Full text available: pdf(1.44 MB) Additional Information: full citation, abstract, citings

On-line analytical processing (OLAP) is a technology that encompasses applications requiring a multidimensional and hierarchical view of data. OLAP applications often require fast response time to complex grouping/aggregation queries on enormous quantities of data. Commercial relational database management systems use mainly multiple one-dimensional indexes to process OLAP queries that restrict multiple dimensions. However, in many cases, multidimensional access methods outperform one-dimensiona ...

17 Tutorial: The relational data model for Design Automation
Mark N. Haynie
June 1983 **Proceedings of the twentieth design automation conference on Design automation**
Full text available: pdf(767.27 KB) Additional Information: full citation, abstract, references, citings, index terms

The relational data model has gained more acceptance in the commercial database environment in recent years. It is now finding its way into the Design Automation (CAD/CAM) area. This tutorial explains what the relational data model is and how database management systems based on it can be used with Design Automation applications.

18 Data access for the masses through OLE DB
José A. Blakeley
June 1996 **ACM SIGMOD Record , Proceedings of the 1996 ACM SIGMOD international conference on Management of data**, Volume 25 Issue 2
Full text available: pdf(1.24 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper presents an overview of OLE DB, a set of interfaces being developed at Microsoft whose goal is to enable applications to have uniform access to data stored in DBMS and non-DBMS information containers. Applications will be able to take advantage of the benefits of database technology without having to transfer data from its place of origin to a DBMS. Our approach consists of defining an open, extensible Collection of interfaces that factor and encapsulate orthogonal, reusable portions ...

19 The Quadtree and Related Hierarchical Data Structures
Hanan Samet
June 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 2
Full text available: pdf(4.87 MB) Additional Information: full citation, references, citings, index terms

20 Industrial sessions: commercial implementation techniques: A compact B-tree
Peter Bumbulis, Ivan T. Bowman
June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data**
Full text available: pdf(825.30 KB) Additional Information: full citation, abstract, references, index terms

In this paper we describe a Patricia tree-based B-tree variant suitable for OLTP. In this variant, each page of the B-tree contains a local Patricia tree instead of the usual sorted array of keys. It has been implemented in iAnywhere ASA Version 8.0. Preliminary experience has shown that these indexes can provide significant space and performance benefits over existing ASA indexes.

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

h be c g e cf c